ISHITA GUPTA

+1(412) 689-5023 | ishitag@cs.cmu.edu | gupta-ish.github.io | Linkedin | Google Scholar

EDUCATION

Carnegie Mellon University — Pittsburgh, PA
Master of Science in Robotics, GPA: 3.83/4.0 — May 2026

- Teaching Assistant (over 400 students): Introduction to Deep Learning, Spring & Fall 2025
- Coursework: Diffusion & Flow Matching, Deep Reinforcement Learning, Generative AI, Learning for 3D Vision

The LNM Institute of Information Technology — Jaipur, India
B.Tech in Computer Science and Engineering — May 2022

- Coursework: Probability & Statistics, Natural Language Processing, Advanced Algorithms, Operating Systems

EXPERIENCE

Nissan Advanced Technology Center | Robotics Research Intern — Santa Clara, CA
Humanoids, VLA Models, Dexterous Manipulation — May 2025 – Aug 2025

- Built **Apple Vision Pro** teleoperation system and curated **800**+ bimanual demonstrations to train the Unitree humanoid for dexterous pick-and-place of warehouse tote boxes **(28 DoF with tactile sensing)**.
- Implemented 2D RGB and **3D point-cloud diffusion policies** for bimanual task learning using the curated dataset.
- Fine-tuned **NVIDIA GR00T N1.5** with LoRA for language-conditioned imitation learning, improving robustness by **3.3x** over baseline diffusion policies on dexterous pick-and-place tasks.
- Developed a scalable humanoid autonomy system with **RL-based whole-body control** (**L4DC '26 Oral**), contributing to the founding of new Humanoid Robotics Lab at Silicon Valley research campus.

Addverb Technologies | Robotics Software Engineer — Noida, India
Software Architecture, Visual SLAM, Computer Graphics — Jan 2022 – Jul 2024

- Implemented the backend of an **ORB-SLAM** system for a quadruped robot, focusing on **pose-graph optimization**, **local bundle adjustment**, and keyframe management to maintain map consistency in GPS-denied environments.
- Engineered real-time, **thread-safe physics simulator** in **modern C++** using OpenGL and NVIDIA PhysX, supporting deterministic **100Hz** control loops and haptic hardware integration.

Google | Software Engineering Intern — Bengaluru, India
Nest Devices, Backend Automation, Python Unit Testing — May 2021 – Aug 2021

- Automated the backend cloud pipeline for camera onboarding, reducing **4-month workflow** to a single execution.
- Built a tool that generated **1,000**+ LOC across multiple languages and automated change-list publishing.

RESEARCH

Robotics Institute, Carnegie Mellon University | Graduate Research Assistant — Pittsburgh, PA
RL for Vision-Language-Action Models (Advised by Prof. Katerina Fragkiadaki) — Dec 2025 – Present

- Developing online RL methods to fine-tune **3D VLAs** conditioned on scene geometry using flow-matching action decoders to improve robustness in bi-manual manipulation and long-horizon planning, evaluated on PerAct2.

PROJECTS

Online Reinforcement Learning for Robotic Foundation Models — CMU

- Fine-tuned OpenVLA-OFT with **GRPO & LoRA**, enabling task adaptation without demonstrations on sparse-reward LIBERO benchmark, boosting task success from **80% to 98%**, preserving **100Hz** control frequency. Report

Humanoid Loco-Manipulation for Warehouse Logistics – *Prof. Guanya Shi* — CMU

- Built full-stack humanoid autonomy system (Unitree G1) for warehouse tote transport with **hierarchical PPO-based whole-body control** and a **FoundationPose + FastSAM** perception pipeline for 6D pose estimation. Website

Semantically Embedded 3D Gaussian Splatting VLAs – *Prof. Shubham Tulsiani* — CMU

- Designed a 3D Gaussian Splatting perception module fused with **NVIDIA GR00T N1.5**, enabling robust occlusion- and height-invariant grasping with a **44%** improvement over vanilla GR00T on Kinova Gen3. Poster

Reinforcement Learning for Autonomous Self-Correction in Language Models — CMU

- Trained LLaMA for autonomous self-correction via a two-stage policy gradient framework with **KL-constrained initialization** and shaped rewards, achieving **57% reduction in answer instability** on MATH500 by mitigating behavior collapse and distribution shift in multi-turn RL. Report

PUBLICATIONS

Zhang, Y., **Gupta, I.**, et al. *FALCON: Force-Adaptive Humanoid Locomotion.* **L4DC 2026, Oral**.
A dual-agent force-adaptive RL framework enabling stable humanoid locomotion under unknown 3D end-effector forces.

SKILLS

Programming & Frameworks: Python, C++, PyTorch, ROS 2, OpenCV, PyTorch3D
Simulation & Tools: MuJoCo, Robosuite, Isaac Sim, Isaac Gym, PhysX, Gazebo, Weights & Biases, Git